arXiv:1606.05242v1 [cs.SI] 16 Jun 2016

Sentiment Aggregate Functions for Political Opinion Polling using Microblog Streams

Pedro Saleiro[1,2,3], Luís Gomes[1,2] and Carlos Soares[1,2,4]
DEI-FEUP[1], Labs Sapo UP [2], LIACC[3], INESC-TEC[4],
University of Porto,
Rua Dr. Roberto Frias, s/n, Porto, Portugal

June 17, 2016

Abstract

The automatic content analysis of mass media in the social sciences has become necessary and possible with the raise of social media and computational power. One particularly promising avenue of research concerns the use of sentiment analysis in microblog streams. However, one of the main challenges consists in aggregating sentiment polarity in a timely fashion that can be fed to the prediction method. We investigated a large set of sentiment aggregate functions and performed a regression analysis using political opinion polls as gold standard. Our dataset contains nearly 233 000 tweets, classified according to their polarity (positive, negative or neutral), regarding the five main Portuguese political leaders during the Portuguese bailout (2011-2014). Results show that different sentiment aggregate functions exhibit different feature importance over time while the error keeps almost unchanged.

Keywords: Sentiment Analysis, Social Media, Political Data Mining

1 Introduction

Surveys and polls using the telephone are widely used to provide information of what people think about parties or political personalities [1]. Surveys randomly select the electorate sample, avoiding selection bias, and are designed to collect the perception of a population regarding some subject, such as in politics or marketing. However this method is expensive and time consuming [1, 2]. Furthermore, over the years it is becoming more difficult to contact people and persuade them to participate in these surveys [3].

On the other hand, online publication of news articles is a standard behavior of news outlets and the raise of social media, namely Twitter and Facebook, has changed the way people interact with news [4, 5]. This way, people are able to react and comment any news in real time. One challenge that several research

works have been trying to solve is to understand how opinions expressed on social media, and their sentiment, can be a leading indicator of public opinion. However, at the same time there might exist simultaneously positive, negative and neutral opinions regarding the same subject. Thus, we need to obtain a value that reflects the general image of each political target in social media, for a given time period. To that end, we use sentiment aggregate functions. In summary, a sentiment aggregate function calculates a global value based on the number of positive, negative, and neutral mentions of each political target, in a given period. We conducted an exhaustive study and collected and implemented several sentiment aggregate functions from the state of the art [4, 6–13].

Thus, the main objective of our work is to study and define a methodology capable of successfully estimating the polls results, based on opinions expressed on social media, represented by sentiment aggregators. We applied this problem to the Portuguese bailout case study, using Tweets from a sample of the Portuguese Tweetosphere and Portuguese polls as gold standard. Given the monthly periodicity of polls, we needed to monthly aggregate data. This approach allows each aggregator value to represent the monthly sentiment for each political party. Due to the absence of a general sentiment aggregate function suitable for different case studies, we decided to include all aggregate functions as features of the regression model. Therefore the learning algorithm is able to adapt to the most informative aggregate functions through time.

In the next Section we review related work. In Section 3 we present the methodology we implemented. We describe data in Section 4 followed by the experimental setup in Section 5. In Section 6 we present and discuss the results we obtained, while Section 7 is reserved for some conclusions taken from our study, and for future work.

2 Related Work

Content analysis of mass media has an established tradition in the social sciences, particularly in the study of effects of media messages, encompassing topics as diverse as those addressed in seminal studies of newspaper editorials [14], media agenda-setting [15], or the uses of political rhetoric [16], among many others. By 1997, Riffe and Freitag [17], reported an increase in the use of content analysis in communication research and suggested that digital text and computerized means for its extraction and analysis would reinforce such trend. Their expectation has been fulfilled: the use of automated content analysis has by now surpassed the use of hand coding [18]. The increase in the digital sources of text, on the one hand, and current advances in computation power and design, on the other, are making this development both necessary and possible, while also raising awareness about the inferential pitfalls involved [19, 20].

One avenue of research that has been explored in recent years concerns the use of social media to predict present and future political events, namely electoral results [4, 6–13]. Although there is no consensus about methods and their consistency [21, 22]. Gayo-Avello [23] summarizes the differences between stud-

ies conducted so far by stating that they vary about period and method of data collection, data cleansing and pre-processing techniques, prediction approach and performance evaluation. One particular challenge when using sentiment is how to aggregate opinions in a timely fashion that can be fed to the prediction method. Two main strategies have been used to predict elections: buzz, i.e., number of tweets mentioning a given candidate or party and the use of sentiment polarity. Different computational approaches have been explored to process sentiment in text, namely machine learning and linguistic based methods [24–26]. In practice, algorithms often combine both strategies.

Johnson et al. [1] concluded that more than predicting elections, social media can be used to gauge sentiment about specific events, such as political news or speeches. Defending the same idea, Diakopoulos el al. [27] studied the global sentiment variation based on Twitter messages of an Obama vs McCain political TV debate while it was still happening. Tumasjan et al. [10] used Twitter data to predict the 2009 Federal Election in Germany. They stated that "the mere number of party mentions accurately reflects the election result". Bermingham et al. [4] correctly predicted the 2011 Irish General Elections also using Twitter data. Gayo-Avello et al. [22] also tested the share of volume as predictor in the 2010 US Senate special election in Massachusetts. On the other hand, several other studies use sentiment as a polls result indicator. Connor et al. [12] used a sentiment aggregate function to study the relationship between the sentiment extracted from Twitter messages and polls results. They defined the sentiment aggregate function as the ratio between the positive and negative messages referring an specific political target. They used the sentiment aggregate function as predictive feature in the regression model, achieving a correlation of 0.80 between the results and the poll results, capturing the important large-scale trends. Bermingham et al. [4] also included in their regression model sentiment features. Bermingham et al. introduced two novel sentiment aggregate functions. For inter-party sentiment, they modified the share of volume function to represent the share of positive and negative volume. For intra-party sentiment , they used a log ratio between the number of positive and negative mentions of a given party. Moreover, they concluded that the inclusion of sentiment features augmented the effectiveness of their model.

Gayo-Avello et al. [22] introduced a different aggregate function. In a two-party race, all negative messages on party $c2$ are interpreted as positive on party $c1$, and vice-versa.

3 Methodology

Figure 1 depicts an overview of the data mining process pipeline applied in this work. To collect and process raw Twitter data, we use an online reputation monitoring platform [28] which can be used by researchers interested in tracking political opinion on the web. It collects tweets from a pre-defined sample of users, applies named entity disambiguation [29] and generates indicators of both frequency of mention and polarity (positivity/negativity) of mentions of entities

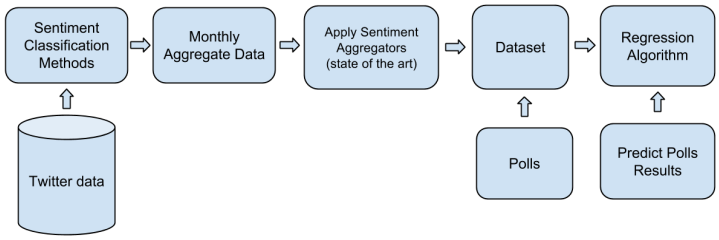

Figure 1: Data mining process pipeline.

across time. In our case, tweets are collected from the stream of 100 thousand different users, representing a sample of the Portuguese community on Twitter.

The platform automatically classifies each tweet according to its sentiment polarity. If a message expresses a positive, negative or neutral opinion regarding an entity (e.g. politicians), it is classified as positive, negative or neutral mention, respectively. The sentiment classifier uses a corpus of 1500 annotated tweets as training set and it is reported an accuracy over 80% using 10-fold cross validation. These 1500 tweets were manually annotated by 3 political science students.

Mentions of entities and respective polarity are aggregated by counting positive, negative, neutral and total mentions for each entity in a given period. Sentiment aggregate functions use these cumulative numbers as input to generate a new value for each specific time period. Since we want to use sentiment aggregate functions as features of a regression model to produce an estimate of the political opinion, we decided to use traditional poll results as gold standard.

3.1 Sentiment Aggregate Functions

The following list presents the sentiment aggregate functions applied to the aggregated data between polls:

- *entity_buzz*: the monthly sum of the number of mentions (buzz) of a given entity (political party leader) between consecutive polls.
- *entity_positives*: the monthly sum of the positively classified mentions of a given entity (political party leader) between consecutive polls.
- *entity_neutrals*: the monthly sum of the neutral classified mentions of a given entity (political party leader) between consecutive polls.
- *entity_negatives*: the monthly sum of the negatively classified mentions of a given entity (political party leader) between consecutive polls.
- *bermingham* [4]: $\log_{10} \frac{entity_posistives+1}{entitty_negatives+1}$

- *berminghamsovn* [4]: $\frac{entity_negatives}{total_negatives}$, $total_negatives$

  corresponds to the sum of the negative mentions of all entities between polls.

- *berminghamsovp* [4]: $\frac{entity_positives}{total_positives}$, $total_positives$ corresponds to the sum of the positives mentions of all entities between polls.

- *connor* [2]: $\frac{entity_positives}{entity_negatives}$

- *gayo* [30]: $\frac{entity_positives+others_negatives}{total_positives+total_negatives}$

- *polarity*: $entity_positives - entity_negatives$

- *polarityONeutral*: $\frac{entity_positives-entity_negatives}{entity_neutrals}$

- *polarityOTotal*: $\frac{entity_positives-entity_negatives}{entity_buzz}$

- *subjOTotal*: $\frac{entity_positives+entity_negatives}{entity}$

- *subjNeuv*: $\frac{entity_positives+entity_negatives}{entity_neutrals}$

- *subjSoV*: $\frac{entity_positives+entity_negatives}{total_positives+total_negatives}$

- *subjVol*: $entity_positives + entity_negatives$

- *share* [4]: $\frac{entity_buzz}{totalbuzz}$

- *shareOfNegDistribution*: $\frac{\frac{entity_negatives}{entity_buzz}}{\sum_{i=0}^{n}\frac{entity_negatives_i}{entity_buzz_i}}$, where n is the number of political entities in the poll

- *normalized_positive*: $\frac{entity_positivesi}{entity_buzz}$

- *normalized_negative*: $\frac{entity_positivesi}{entity_buzz}$

- *normalized_neutral*: $\frac{entity_positivesi}{entity_buzz}$

- *normalized_bermingham*: $\log_{10}\frac{normalized_positives+1}{normalized_negatives+1}$

- *normalized_connor*: $\frac{normalized_positives}{normalized_negatives}$

- *normalized_gayo*:

  $$\frac{normalized_positives+normalized_others_negatives}{normalized_total_positives+normalized_total_negatives}$$

- *normalized_polarity*:

  $$normalized_positives - normalized_negatives$$

The sentiment aggregate functions are used as features in the regression models.

Table 1: Distribution of positive, negative and neutral mentions per political party

	Negative	Positive	Neutral	Total Mentions
PSD	69 723	121	37 133	106 977
PS	28 660	225	15 326	44 211
CDS	41 935	51	17 554	59 540
CDU	2 445	79	5 604	8 128
BE	9 603	306	4 214	14 123

4 Data

The data used in this work consists of tweets mentioning Portuguese political party leaders and polls from August 2011 to December 2013. This period corresponds to the Portuguese bailout when several austerity measures were adopted by the incumbent right wing governmental coalition of PSD and CDS.

4.1 Twitter

The Twitter data set contains 232 979 classified messages, collected from a network of 100 thousand different users classified as Portuguese. Table 1 presents the distribution of positive, negative, and neutral mentions of the political leaders of the 5 most voted political parties in Portugal (PSD, PS, CDS, PCP and BE). The negative mentions represent the majority of the total mentions, except for CDU where the number of negative mentions is smaller than the neutral ones. The positive mentions represent less than 1% of the total mentions of each party, except for BE where they represent 2% of the total mentions. The most mentioned parties are PS, PSD and CDS. The total mentions of these three parties represent 90% of the data sample total mentions.

Figure 2 depicts the time series of the *berminghamsovn* (negatives share) sentiment aggregate function. The higher the value of the function the higher is the percentage of negative tweets mention a given political entity in comparison with the other entities. As expected, Pedro Passos Coelho (PSD) as prime-minister is the leader with the higher score through the all time period under study. Paulo Portas (CDS) leader of the other party of the coalition, and also member of the government is the second most negatively mentioned in the period, while António José Seguro (PS) is in some periods the second higher. PSD and CDS are the incumbent parties while PS is the main opposition party in the time frame under study. PSD and CDS as government parties were raising taxes and cutting salaries. PS was the incumbent government during the years that led to the bailout and a fraction of the population considered responsible for the financial crisis. The bailout and the consequent austerity measures could explain the overwhelming percentage of negative mentions although we verified that in other time periods the high percentage of negatives mentions remains. We can say that Twitter users of this sample when mentioning political leaders

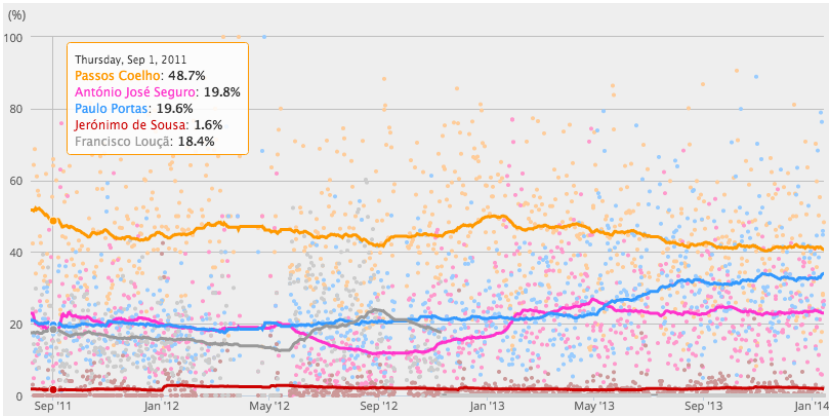

Figure 2: Negatives share (*berminghamsovn*) of political leaders in Twitter.

on their tweets tend to criticize them.

4.2 Political Opinion Polls

The polling was performed by Eurosondagem, a Portuguese private company which collects public opinion. This data set contains the monthly polls results of the five main Portuguese parties, from June 2011 to December 2013. Figure 3 represents the evolution of Portuguese polls results. We can see two main party groups: The first group, where both PSD and PS are included, has a higher value of vote intention (above 23%). PSD despite starting as the preferred party in vote intention, has a downtrend along the time, losing the leadership for PS in September 2012. On the other hand, PS has in general an uptrend. The second group, composed by CDS, PCP and BE, has a vote intention range from 5% to 15%. While CDS has a downtrend in public opinion, PCP has an ascendent one. Although the constant tendencies (up- and downtrends), we noticed that the maximum variation observed between two consecutive months is 3%. In June 2013 there was political crises in the government when CDS threaten to leave the government coalition due to the austerity measures being implemented and corresponds to the moment when PS takes the lead in the polls.

5 Experimental Setup

We defined the period of 2011 to December 2012 as training set and the all year of 2013 as test set. We applied a sliding window setting in which we start to predict the poll results of January 2013 using the previous 16 months as training set. The second poll to estimate is February 2013 and we train a new model using the previous 16 months to the target month under prediction.

- Training set – containing the monthly values of the aggregators (both

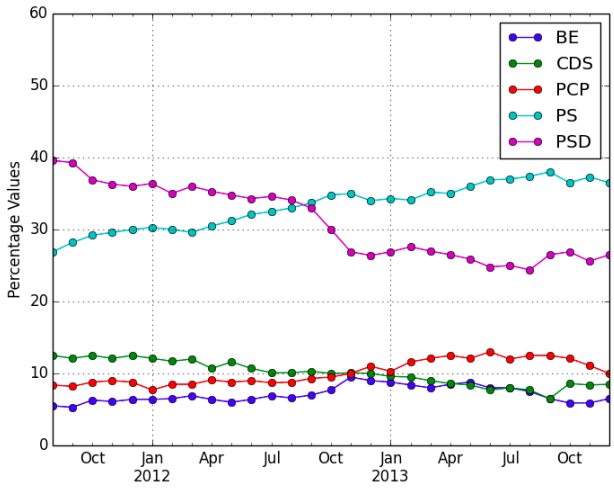

Figure 3: Representation of the monthly poll results of each political candidate

sentiment and buzz aggregator) for 16 months prior the month intended to be predicted.

- Testing set - containing the values of the aggregators (both sentiment and buzz aggregator) of the month intended to be predicted.

1. We select the values of the aggregators of the 16 months prior January 2013 (September 2011 to December 2012).
2. We use that data to train our regression model.
3. Then we input the aggregators' values of January 2013 - the first record of the testing set - in the the trained model, to obtain the poll results prediction.
4. We select the next month of the testing set and repeat the process until all months are predicted.

The models are created using two regression algorithms: a linear regression algorithm (Ordinary Least Squares - OLS) and a non-linear regression algorithm (Random Forests - RF). We also run an experiment using the derivative of the polls time series as gold standard, i.e., poll results variations from poll to poll. Thus, we also calculate the variations of the aggregate functions from month to month as features. Furthermore, we repeat each experiment including and excluding the lagged self of the polls, i.e., the last result of the poll for a given candidate (y_{t-1}) or the last polls result variation (Δy_{t-1}) when predicting polls variations. We use Mean Absolute Error (MAE) [31] as evaluation measure, to determine the absolute error of each prediction. Then, we calculate the average of the twelve MAE's so we could know the global prediction error of our model.

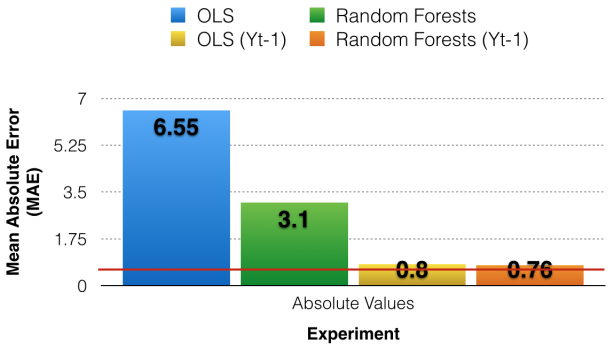

Figure 4: Error predictions for polls results.

$$MAE = \frac{\sum_{i=1}^{n} |f_i - y_i|}{n} \tag{1}$$

n is the number of forecasts, f_i is the model's forecast and y_i the real outcome.

6 Results and Discussion

In this Section we explain in detail the experiments and its results. We perform two different experiments: (1) using absolute values and (2) using monthly variations.

6.1 Predicting Polls Results

In this experiment, the sentiment aggregators take absolute values in order to predict the absolute values of polls results. Mathematically speaking, this experiment can be seen as: $y \leftarrow \{y_{t-1}, buzzAggregators, sentimentAggregators\}$. In figure 4 we see the global errors we obtained.

The results shows that we obtain an Mean Absolute Error for the 5 parties poll results over 12 months of 6.55 % using Ordinary Least Squares and 3.1 % using Random Forests. The lagged self of the polls, i.e., assuming the last known poll result as prediction results in a MAE of 0.61 % which was expectable since the polls exhibit slight changes from month to month. This experiment shows that the inclusion of the lagged self (y_{t-1}) produces average errors similar to the lagged self.

6.2 Predicting Polls Results Variation

According to our exploratory data analysis, the polls results have a small variation between two consecutive months. Thus, instead of predicting the abso-

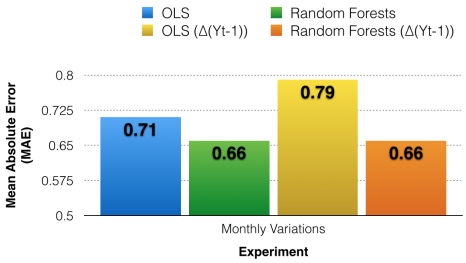

Figure 5: Error predictions for polls results variation.

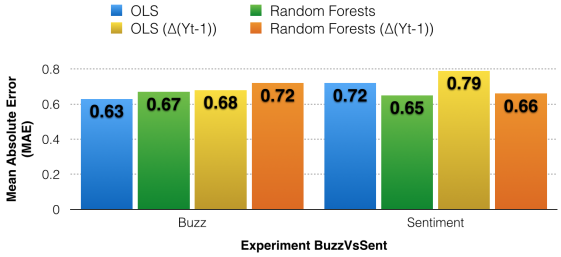

Figure 6: Mean absolute error buzz vs sentiment.

lute value of poll results, we tried to predict the variation, $\Delta y \leftarrow \{\Delta(y_{t-1}), \Delta buzzAggregators, \Delta sentimentAggregators\}$

In this particular experiment, the inclusion of the Δy_{t-1} as feature in the regression model has not a determinant role (figure 5). Including that feature we could not obtain lower MAE than excluding it. It means that the real monthly poll variation is not constant over the year. In general, using a non-linear regression algorithm we obtain lower MAE. The results show that when leading with polls results with slight changes from poll to poll it makes sense to transform the dataset by derivation.

6.2.1 Buzz and Sentiment

Several studies state that the buzz has predictive power and reflects correctly the public opinion on social media. Following that premise, we trained our models with buzz and sentiment aggregators separately to predict polls variations:

- $\Delta y \leftarrow \{\Delta(y_{t-1}), \Delta buzzAggregators\}$
- $\Delta y \leftarrow \{\Delta(y_{t-1}), \Delta sentimentAggregators\}$

This experiment allowed us to compare the behavior of buzz and sentiment aggregators.

According to figure 6, buzz and sentiment aggregators have similar results. Although the OLS algorithm combined only with buzz aggregators has a slightly

lower error than the other models, it is not a significant improvement. These results also show that Random Forests algorithm performs the best when combined only with sentiment aggregators.

6.3 Feature Selection

One of the main goals of our work is to understand which aggregator (or group of aggregators) better suits our case study. According to the previous experiments, we can achieve lower prediction errors when training our model with buzz and sentiment aggregators separately. However, when training our model with these two kinds of aggregators separately, we are implicitly performing feature selection. We only have two buzz features (*share* and *total_mentions*). Due to that small amount of features, it was not necessary to perform any feature selection technique within buzz features. Thus, we decided to apply a feature selection technique to the sentiment aggregators, in order to select the most informative ones to predict the monthly polls results variation. We use univariate feature selection, selecting 10% of the sentiment features (total of 3 features). Using this technique, the Random Forests' global error raise from 0.65 to 0.73. However, OLS presents an MAE drop from 0.72 to 0.67. Another important fact to notice is that if we perform univariate feature selection to all aggregators (buzz and sentiment), we will achieve the same MAE value that when applied only to sentiment aggregators.

We try a different approach and perform a recursive feature elimination technique. In this technique, features are eliminated recursively according to a initial score given by the external estimator. This method allow us to determine the number of features to select. Thus, also selecting 3 features, the OLS' MAE drop to 0.63. Once again, none of the buzz features were selected. Furthermore, both feature selection techniques select different features for each monthly prediction.

6.4 Feature Importance

We select the Random Forest model of monthly variations to study the features importance as depicted in figure 7. The higher, the more important the feature. The importance of a feature is computed as the (normalized) total reduction of the criterion brought by that feature. It is also known as the Gini importance. Values correspond to the average of the Gini importance over the different models trained in the experiments. The single most important feature is the *bermingham* aggregate function, followed by *neutrals*. It is important to notice that when combining all the aggregate functions as features in a single regression model, the *buzz* does not comprises a high Gini importance, even if when used as a single feature produces similar results with the sentiment aggregate functions. In general, the standard deviation of the GIni importance is relatively high. This has to due with our experimental setup, as the values depicted in the bar chart correspond to the average of the Gini importance over 12 different models (12 months of testing set). Therefore, feature importances

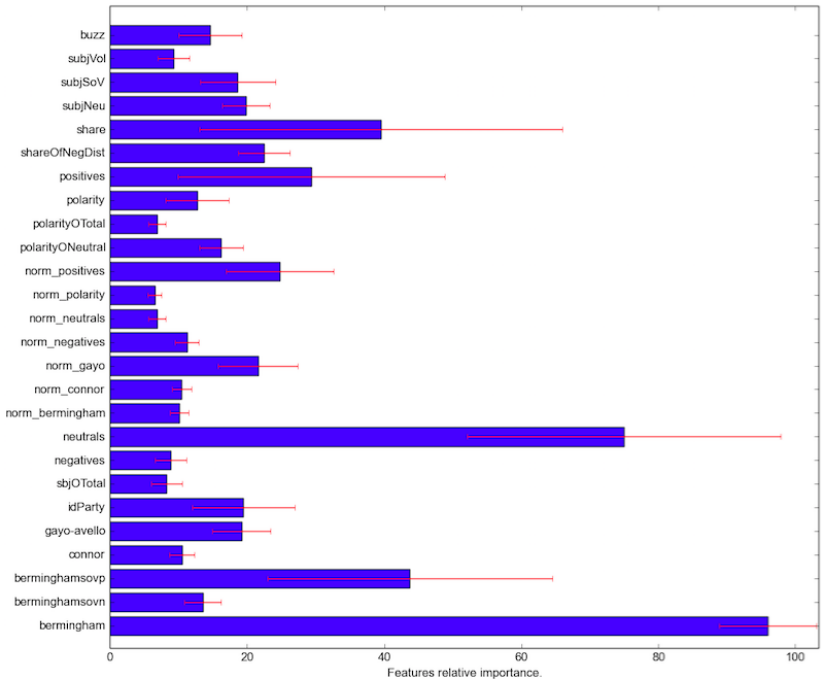

Figure 7: Aggregate functions importance in the Random Forests models.

vary over time while the MAE tends to remain unchanged. We can say that different features have different informative value over time and consequently it is useful to combine all the sentiment aggregation functions as features of the regression models over time.

7 Conclusions

We studied a large set of sentiment aggregate functions to use as features in a regression model to predict political opinion poll results. The results show that we can estimate the polls results with low prediction error, using sentiment and buzz aggregators based on the opinions expressed on social media. We introduced a strong baseline for comparison, the lagged self of the polls. In our study, we built a model where we achieve the lowest MAE using the linear algorithm (OLS), combined only with buzz aggregators, using monthly variations. The model has an MAE of 0.63%. We performed two feature selection techniques: (1) Univariate feature selection and (2) recursive feature elimination. Applying the recursive technique to the sentiment features, we can achieve an MAE of 0.63, equating our best model. Furthermore, the chosen features are not the same in every prediction. Regarding feature importance analysis our experiments showed that *bermingham* aggregate function represents the higher

Gini importance in the Random Forests model. The next immediate step is to implement a methodology using time series analysis. Furthermore, it is desirable to test this methodology with difference data sources, such as Facebook messages, blogs or news. Other alternative approach we intend to implement and evaluate is to interpret this problem as a classification problem - predict only the changing direction of opinion poll result (i.e., up or down).

References

[1] C. Johnson, P. Shukla, and S. Shukla. On classifying the political sentiment of tweets. 2012.

[2] B. Connor, R. Balasubramanyan, B. Routledge, and N. Smith. From tweets to polls : Linking text sentiment to public opinion time series. In *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*, pages 122–129, May 2010.

[3] A. Kohut, S. Keeter, C. Doherty, M. Dimock, A. Directors, and L. Christian. Assessing the representativeness of public opinion surveys. 2012.

[4] A. Bermingham and A. Smeaton. On using twitter to monitor political sentiment and predict election results. *Workshop at the International Joint Conference for Natural Language Processing (IJCNLP)*, November 2011.

[5] Pedro Saleiro, Jorge Teixeira, Carlos Soares, and Eugénio Oliveira. Timemachine: Entity-centric search and visualization of news archives. In *Advances in Information Retrieval: 38th European Conference on IR Research, ECIR 2016, Padua, Italy, March 20-23, 2016. Proceedings*, pages 845–848. Springer International Publishing, 2016.

[6] Andranik Tumasjan, Timm Oliver Sprenger, Philipp G Sandner, and Isabell M Welpe. Predicting elections with twitter: What 140 characters reveal about political sentiment. *ICWSM*, 10, 2010.

[7] Micol Marchetti-Bowick and Nathanael Chambers. Learning for microblogs with distant supervision: Political forecasting with twitter. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, EACL '12. Association for Computational Linguistics, 2012.

[8] Pawel Sobkowicz, Michael Kaschesky, and Guillaume Bouchard. Opinion mining in social media: Modeling, simulating, and forecasting political opinions in the web. *Government Information Quarterly*, 29(4):470 – 479, 2012. Social Media in Government - Selections from the 12th Annual International Conference on Digital Government Research.

[9] Avishay Livne, Matthew P Simmons, Eytan Adar, and Lada A Adamic. The party is over here: Structure and content in the 2010 election. In *ICWSM*, 2011.

[10] Andranik Tumasjan, Timm O Sprenger, Philipp G Sandner, and Isabell M Welpe. Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proceedings of the fourth international AAAI conference on weblogs and social media*, 2010.

[11] D Gayo-Avello. I wanted to predict elections with twitter and all i got was this lousy paper a balanced survey on election prediction using twitter data. *arXiv preprint arXiv:1204.6441*, 2012.

[12] Brendan O'Connor, Ramnath Balasubramanyan, Bryan R Routledge, and Noah A Smith. From tweets to polls: Linking text sentiment to public opinion time series. In *Proceedings of the International AAAI Conference on Weblogs and Social Media*, 2010.

[13] Jessica Chung and Eni Mustafaraj. Can collective sentiment expressed on twitter predict political elections. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence. San Francisco, CA, USA*, 2011.

[14] Harold Dwight Lasswell. *The comparative study of symbols: An introduction*. Number 1. Stanford University Press, 1952.

[15] Maxwell E McCombs and Donald L Shaw. The agenda-setting function of mass media. *Public opinion quarterly*, 36(2), 1972.

[16] Matthew C Moen. Ronald reagan and the social issues: Rhetorical support for the christian right. *The Social Science Journal*, 27(2):199–207, 1990.

[17] Daniel Riffe and Alan Freitag. A content analysis of content analyses: Twenty-five years of journalism quarterly. *Journalism & Mass Communication Quarterly*, 74(3), 1997.

[18] Kimberly A Neuendorf. *The content analysis guidebook*. Sage, 2002.

[19] Daniel J Hopkins and Gary King. A method of automated nonparametric content analysis for social science. *American Journal of Political Science*, 54(1), 2010.

[20] Justin Grimmer and Brandon M. Stewart. Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 2013.

[21] Panagiotis T. Metaxas, Eni Mustafaraj, and Dani Gayo-Avello. How (Not) to Predict Elections. *2011 IEEE Third Int'l Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third Int'l Conference on Social Computing*, October 2011.

[22] Daniel Gayo Avello, Panagiotis T Metaxas, and Eni Mustafaraj. Limits of electoral predictions using twitter. In *Proceedings of the International Conference on Weblogs and Social Media*, 2011.

[23] Daniel Gayo-Avello. A meta-analysis of state-of-the-art electoral prediction from twitter data. *Social Science Computer Review*, page 0894439313493979, 2013.

[24] Bo Pang and Lillian Lee. Opinion mining and sentiment analysis. *Found. Trends Inf. Retr.*, 2(1-2), 2008.

[25] Efthymios Kouloumpis, Theresa Wilson, and Johanna Moore. Twitter sentiment analysis: The good the bad and the omg. In *Proceedings of the International Conference on Weblogs and Social Media*, 2011.

[26] Preslav Nakov, Sara Rosenthal, Zornitsa Kozareva, Veselin Stoyanov, Alan Ritter, and Theresa Wilson. Semeval-2013 task 2: Sentiment analysis in twitter. In *Proceedings of the International Workshop on Semantic Evaluation (SemEval 2013)*, 2013.

[27] N. Diakopoulos and D. Shamma. Characterizing debate performance via aggregated twitter sentimen. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI'10*, pages 1195–1198. ACM Press, April 2010.

[28] Pedro Saleiro, Silvio Amir, Mário Silva, and Carlos Soares. Popmine: Tracking political opinion on the web. In *Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing (CIT/IUCC/DASC/PICOM), 2015 IEEE International Conference on*, pages 1521–1526. IEEE, 2015.

[29] Pedro Saleiro, Luıs Rei, Arian Pasquali, and Carlos Soares. Popstar at replab 2013: Name ambiguity resolution on twitter. In *CLEF 2013 Eval. Labs and Workshop Online Working Notes*, 2013.

[30] P. Metaxas, E. Mustafaraj, and D. Gayo-Avello. How (not) to predict elections. *2011 IEEE Third Int'l Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third Int'l Conference on Social Computing*, pages 165–171, October 2011.

[31] J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers, 2006.